

Express Mailing Label No.: ET588572945US

PATENT APPLICATION

Docket No.: 1200.2.34

IBM Docket No.: SJO920010113US1

UNITED STATES PATENT APPLICATION

of

Anthony Harty

Theresa Segura

and

Brian McKean

for

**HIERARCHICAL MULTI-COMPONENT TRACE FACILITY
USING MULTIPLE BUFFERS PER COMPONENT**

2009-07-20 14:00:00

1 **HIERARCHICAL MULTI-COMPONENT TRACE FACILITY**
2 **USING MULTIPLE BUFFERS PER COMPONENT**

3 **BACKGROUND OF THE INVENTION**

4 **1. The Field of the Invention**

5 The invention relates to embedded systems used to locate and document program
6 errors. More specifically, the present invention is related to systems and methods for
7 implementing multiple trace buffers to identify errors within a program.

8 **2. The Relevant Art**

9 Knowledge of the internal operation of code and micro-code within a computer is
10 useful for debugging, optimization, and design verification. Viewing the internal operation
11 of programs, known as tracing, provides a view into the behavior of the programs by
12 recording the time and details of the state of the programs at relevant points in the operation
13 of the programs.

14 Tracing is a broadcast form of inter-process communication with many source
15 processes (e.g., writer entities) and sink processes (e.g., reader entities) capable of observing
16 each other's execution. A trace, therefore, consists of embedded code that chronicles the
17 actions and results of program execution. Specifically, the trace provides a detailed record of
18 the program's execution path by, for instance, placing an application program under
19 observation by a dedicated routine that monitors the progress of the program.

20 Performance information obtained by writer entities engaged in tracing operations is
21 typically provided to the registry in the form of trace scripts. Each trace script has a variable
22 length and is stored in a portion of the registry configured to accommodate such variable-
23 length messages. This storage location is generally known as a trace buffer.

24 A plurality of messages and corresponding fragments are generally interleaved
25 within the trace buffer, which is typically shared among the software components or other
26 entities being monitored. The trace buffer may be accessed by multiple writer entities

1 attempting to load messages into the buffer and multiple reader entities attempting to retrieve
2 those messages.

3 Traces provide computer engineers with a view of process and data states while the
4 computer systems are operating in real-time. Hardware engineers often use traces to
5 determine how new computer hardware architectures will perform with many different types
6 of operating systems and application programs. Specific designs of hardware structures,
7 such as instruction processors and data memories, can have drastically different and
8 sometimes unpredictable utilizations for the same sets of instructions and data. It is
9 important that any flaws in the hardware design are found before the design is finalized.

10 Software engineers are also often required to identify critical information about code
11 segments and data structures. For example, it is useful for compiler writers to know how the
12 compiler schedules instructions for execution and how well conditional branches are
13 predicted to provide input for code optimization. It is similarly useful for software engineers
14 debugging programs to investigate the exact execution and data flow in an errant code
15 segment.

16 One manner of providing a trace system to locate errors in a coded program involves
17 embedding the trace system within the program it is to debug. An embedded trace system
18 allows the system to evaluate each section of the code while the computer system is running
19 various programs. By embedding the system, the user is not required to start the debugging
20 process, provide commands or prompts, or wait while the tracing system scans the program.

21 An embedded tracing system also decreases the possibility of erroneous changes
22 being made to the tracing system. Typically in an embedded trace system, users are unable
23 to access the code of the tracing system, change the code, and render the tracing system
24 unmanageable. This provides for a certain level of security to the code containing the
25 tracing system.

26 As it is known in the art, computer systems generally include a central processing
unit, a memory subsystem, and a storage subsystem. According to a networked or enterprise

As the tracing commences within prior art systems, both messages and events are stored in the common buffer. Due to the size constraint, the buffer frequently becomes full and incapable of storing additional information. Previous messages and events that

1 had been stored are either discarded or overwritten to provide room for new messages and
2 events within the buffer. The user is thus unaware of the earlier events, which might be
3 pertinent and even critical to the debugging or other operation on the system being traced.

4 From the above discussion, it can be seen that it would be a beneficial addition in
5 the art to provide an improved tracing technique. Particularly, it would be advantageous
6 to provide an improved trace buffer that is more flexible and less subject to information
7 overruns.

8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26

1 **OBJECTS AND BRIEF SUMMARY OF THE INVENTION**

2 The tracing facility of the present invention has been developed in response to the
3 present state of the art, and in particular, in response to the problems and needs in the art
4 that have not yet been fully solved by currently available tracing facilities. Accordingly,
5 it is an overall object of the present invention to provide a tracing facility that overcomes
6 many or all of the above-discussed shortcomings in the art.

7 The tracing facility of the present invention is highly suitable for locating errors in
8 an embedded system. The embedded system in one embodiment comprises a segment of
9 program code, a plurality of buffers, and a tracing module configured to locate events
10 within the segment of program code which are useful for finding and correcting errors.
11 The tracing module is also configured to selectively transmit the events to the plurality of
12 buffers.

13 Each buffer in the plurality of buffers may be configured to store a different type
14 of event from the other buffers. The different types of events stored within the plurality
15 of buffers may comprise, by way of example, errors, warnings, and messages. The
16 plurality of buffers may be assigned to a single segment of program code or may be
17 distributed among a plurality of functional components of a computer program, with each
18 buffer or set of buffers assigned to a different functional component of the program code.

19 A merging module within or external to the embedded system of the tracing
20 facility is preferably configured to combine the events from the plurality of buffers into a
21 common list of events. The list of events may be organized in a chronological order
22 according to a timestamp placed on each event stored in a buffer. The merging module is
23 preferably also configured to selectively combine the events. Accordingly, each event
24 stored in the plurality of buffers may be stored together with information about the event
25 such as the type of event, the location of the program code where the event took place,
26 and the time of the event.

1 A method of the present invention for providing an embedded system with a trace
2 facility is also provided. The method comprises providing a plurality of buffers
3 configured to store events useful in finding and correcting errors, tracing the events
4 within a segment of program code, and selectively storing the events within the plurality
5 of buffers. The events are preferably stored according to the types of the events,
6 including errors, warnings, and messages.

7 The method also preferably comprises distributing the plurality of buffers among
8 a plurality of functional components of the program code, with each buffer assigned to a
9 different functional component. The events may then be merged from the plurality of
10 buffers into a common list of events, organized chronologically. A timestamp is
11 preferably placed on each event stored in the plurality of buffers to facilitate
12 chronological organization.

13 These and other objects, features, and advantages of the present invention will
14 become more fully apparent from the following description and appended claims, or may
15 be learned by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

In order that the manner in which the advantages and objects of the invention are obtained will be readily understood, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments thereof which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

Figure 1 is a block diagram of a computer system suitable for implementing the present invention;

Figure 2 is a schematic block diagram illustrating one embodiment of a trace facility of the present invention;

Figure 3 is a schematic flow chart diagram illustrating one embodiment of a process of the present invention for developing a software product using a tracing program code;

Figure 4 is a schematic flow chart diagram illustrating one embodiment of a method for tracing program code within a computer system; and

Figure 5 is a schematic flow chart diagram illustrating one embodiment of a method for merging events stored in separate buffers in accordance with the present invention.

BRIAN C. KUNZLER
ATTORNEY AT LAW
10 WEST 100 SOUTH, SUITE 425
SALT LAKE CITY, UTAH 84101

Parameter	Value	Unit	Parameter	Value	Unit
α	0.001		β	0.001	
γ	0.001		δ	0.001	
ϵ	0.001		ζ	0.001	
η	0.001		θ	0.001	
ι	0.001		κ	0.001	
λ	0.001		μ	0.001	
ν	0.001		ξ	0.001	
\omicron	0.001		π	0.001	
ρ	0.001		σ	0.001	
τ	0.001		υ	0.001	
ϕ	0.001		χ	0.001	
ψ	0.001		ω	0.001	
Ω	0.001		Θ	0.001	
Φ	0.001		Ψ	0.001	
Υ	0.001		Ξ	0.001	
Λ	0.001		Σ	0.001	
Π	0.001		Υ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	
Σ	0.001		Γ	0.001	
Δ	0.001		Σ	0.001	
Γ	0.001		Δ	0.001	

13
14
15
16
17

18
19
20
21
22

23

24

25

26

27

1 In the depicted embodiment, the trace facility 200 is implemented to chronicle or
2 document errors in the program 234 and may be embedded within the program 234. Within
3 the trace facility 200, an input module 202 is configured to receive input information from
4 the user to the system. A keyboard 124 or mouse 126 of Fig. 1 may comprise suitable
5 devices for communicating with the input module 202. An output module 204 is configured
6 to output information from the system to the user. Components interconnected to a display
7 adapter 114 such as a computer monitor or a printer may be used to communicate with the
8 output module 204.

9 The program 234 may comprise any type of computer program within which it is
10 desirable to locate errors as they occur in real time. In one embodiment, the program 234 is
11 broken into logical components. Shown by way of example for illustration purposes only are
12 components including an initialize component 206, a read component 208, and a write
13 component 210. Each component of the program 234 may be configured to have its code
14 traced by a tracing module 212. The tracing module 212 is thus configured to record events
15 from the operation of selected segments of code from the individual components that
16 together comprise the system program code. Records of these events, which in the prior art
17 are stored in a single trace buffer, are stored in a plurality of trace buffers under the present
18 invention.

19 In one embodiment, the tracing module 212 is a function that is called with
20 parameters indicating the event that has occurred. The tracing module 212 is preferably
21 placed throughout the program code being traced on an as-needed basis. The tracing module
22 212 thus is a valuable tool that records events for later analysis to facilitate the location of
23 errors.

24 In the depicted embodiment, nine trace buffers 214 are depicted. Nevertheless, any
25 suitable number of trace buffers may be used. Also in the depicted embodiment, the trace
26 buffers 214 are categorized in different manners. For instance, the plurality of trace buffers
27 214 may be organized by the type of event stored. In the depicted embodiment, three types

1 of event buffers -- errors, warnings, and messages -- are shown. By separating the buffers
2 into types, certain problems are alleviated. For instance, there may be a hierarchy among the
3 types, so more important events such as errors get preferential treatment. Also, using this
4 scheme, sufficient space can be allocated to the more important events so that, for instance
5 large amounts of warnings do not cause errors to overflow.

6 Another manner of categorizing the plurality of buffers is by functional components
7 of the program 234. For instance, the depicted program 234 is shown by way of example to
8 be divided conceptually (although the code may also be physically divided such as by
9 procedures, objects, or data structures) into three components, the initialize components 206,
10 the read component 208, and the write component 210. Three buffers are shown allocated to
11 each component.

12 Accordingly, during operation, when the tracing module 212 records errors in the
13 code of the initialize component 206, the errors are stored as initialize component errors 214.
14 Events occurring within the initialize component 206 that are traced by the tracing module
15 212 and are evaluated as warnings are stored as an initialize warning 216, and events traced
16 as messages are stored as an initialize message 218. Events located by the tracing module
17 212 as occurring within the read component 208 and write component 210 are handled in a
18 similar manner according to the assigned buffers.

19 After the code for each of the individual components of the system program code
20 234 has been traced and evaluated by the tracing module 212, and a user is ready to view the
21 results, a merging module 232 may be used to combine the stored codes from the disparate
22 buffers 214. The merging module 232 may, for instance, be configured to organize all the
23 error events, including, the errors found in the initialize component 206, the errors found in
24 the read component 208, and the write component 210 into a single list. Similarly, all
25 warnings may be combined, and all messages may be combined. Furthermore, all events for
26 a component may be combined, and all events for the program 234 may be combined. The
27 events may be combined in any suitable combination. This allows the user, or system

1 administrator to identify errors, warnings, and messages found in each component in a short
2 amount of time. Preferably, each event is time stamped when placed in the buffers 214 so
3 that when the events are combined they may be listed in chronological order.

4 Figure 3 is a schematic flow chart diagram illustrating one embodiment of a process
5 300 for developing a software product using a tracing facility of the present invention. While
6 the method 300 may be conducted independently of the system of Figures 1 and 2, the
7 method 300 will be described herein by way of example with reference to the systems 100
8 and 200 of Figures 1 and 2. The process starts 301, and a program is provided 302. A trace
9 facility 200 is then provided 304, and in one embodiment is embedded 302 in the program
10 234.

11 During development of the software, the developers perform tests on the software to
12 thoroughly find and correct any errors (bugs) located within the software. This step is
13 commonly known as debugging 306 the product. The tracing facility may be used at this
14 step in a manner that is described in greater detail with reference to Figure 4. Once the
15 product has been thoroughly tested and debugged, the product is shipped 308. The product is
16 then received and begins operation 310. In one embodiment, the tracing facility is embedded
17 within the program 234 and may be transparent to the user.

18 The developer or the user may use the trace facility to trace 312 and continue
19 debugging the product at any time during the life of the product on-site and in real time. In
20 so doing, events discovered within the program 234 are optionally merged into a single
21 buffer 313. Thus, in one embodiment, using the trace facility errors, warnings, and messages
22 are detected within the product on-site 314 and in real time. This allows the program code to
23 be easily tested and debugged even after the product has been shipped and installed. Located
24 errors, warnings, and messages are then analyzed 326.

25 Errors and issues arising from warnings may then be corrected in one embodiment
26 by recoding 328 the program. Once the program has been recoded, tests are conducted to
27 determine whether the program is operable 330 or otherwise suited to its purpose. If so, the

1 method 300 ends 318. If more testing is needed, either because the program is visibly
2 unsatisfactory or because errors, warnings, and/or messages are still present, the method 300
3 again returns to trace 312 the program code and locate events in search of a solution to the
4 indicated errors and issues.

5 Figure 4 is a flow chart diagram illustrating one embodiment of a method 400 for
6 tracing a segment of program code in accordance with the present invention. The method
7 400 may be employed in one embodiment as the "Trace Program Code" step 312 of Figure 3.

8 The method 400 begins 402, and the tracing facility is invoked 404. The tracing
9 facility is then used to trace 406 the program code of each component 206, 208, 210 within
10 the program 234. As the program code in a component is traced, events occur 408. Each
11 traced event is noted and is preferably time stamped 410. Time stamping events 410 permits
12 post-processing of the buffers, in such a way that events from all buffers or selected buffers
13 may be merged chronologically in a manner to be discussed. When an event is identified as
14 an error, a record of the event is made and then is time stamped 410 and stored 412 in an
15 error buffer. The method 400 then returns to step 404 and continues checking for more
16 events.

17 Referring now to step 414, a recognized event may be identified as a warning.
18 A warning generally comprises code that may potentially become an error and render the
19 program code of the system inoperable or unsuited for its intended purpose. Following the
20 discovery of a warning event 414, a copy of the event is time stamped 410 and stored 416 in
21 a separate buffer configured to store warning events of the particular component being
22 traced.

23 Referring to step 418 an event may also be recognized as a message event.
24 Messages allow the user or system administrator to review past history of the traces and to
25 verify the function of the individual components program code. A copy of the message event
26 is time stamped 410 and then stored in a separate buffer 420 configured to store message
27 events of the particular component. Three types of events are illustrated in Figure 3.

1 Nevertheless, it should be readily apparent that events may be tracked that differ from those
2 shown.

3 After the event is stored in buffers 416, 418, or 420, the method 400 continues
4 tracing the program code until a terminate command is received 422. The terminate
5 command may be generated by a user or as a consequence of the termination of the program
6 being traced. When a terminate command is received 422, the method 400 ends.

7 Table 1 is one example of the possible contents of event buffers 214 in which events
8 are recorded in accordance with the methods 300, 400.

9 **Component 1 Trace**

10 **Errors**

11	Time	Function	State1	State2	Vars1	Vars2	Event	Free1
12	25684779	FunctionName1	0000	0000	0005	0151	00AC	0000
13	25784300	FunctionName2	9B24	0000	0001	0005	013F	0000

14 **Warnings**

15	Time	Function	State1	State2	Vars1	Vars2	Event	Free1
16	25784901	FunctionName2	9B23	0000	0005	0151	013A	0000
17	25790200	FunctionName3	9B24	0000	0001	0005	0122	0000
18	25790999	FunctionName3	9B24	0000	0001	0005	0122	0000

19 **Messages**

20	Time	Function	State1	State2	Vars1	Vars2	Event	Free1
21	21684712	FunctionName4	9B24	0000	0005	0151	010F	0000
22	22734407	FunctionName2	9B24	0000	0001	0005	01F5	0000
23	23334000	FunctionName1	9B24	0000	0000	000F	01F3	0000
24	23785393	FunctionName2	9B24	0000	0001	0005	01F5	0000

25 **Table 1**

26 Table 1 contains several columns with headings and data or names below each
27 heading. The first column labeled time, indicates the time stamp of the event. The time

1 stamp allows the user to know when the event was recorded by the tracing facility. The
2 measurement of the timestamp is shown in the depicted embodiment as clock cycles of
3 the CPU.

4 Program code is typically organized with different functions. The second
5 column labeled function, indicates the function name within the code where an event was
6 discovered.

7 The third and fourth columns labeled State1 and State2, are examples of a certain
8 state that a functional area or component might be in at that point in time. The state is
9 usually a global variable which implies that the function might be left for a certain
10 amount of time, and upon returning to that function the state it was in is stored within the
11 code. An example of State1 might be a read state, and an example of State2 might be a
12 write state. The code listed below State1 and State2 is a pre-defined position of that state.
13 For example, the code "0000" below State1 in table 1 under Errors, might indicate a
14 read-ahead state, while "9B24" might indicate a read-backwards state.

15 The next columns in table 1 labeled Vars1 and Vars2, indicate the effected
16 variable or object within the function where an event was discovered. The values listed
17 below Vars1 and Vars2 are examples of actual values of the variable or object. The value
18 "0005" listed in Table 1 under Errors and Vars1 may indicate the value of Vars1 at that
19 point in time. The user is then able to determine the cause of the error, warning, or
20 message. For example, the value may be too large or too small for the parameters of that
21 particular function within the code being traced. Each function has several variables that
22 may be listed and their values recorded.

23 The next column labeled Event, indicates a pre-programmed event that needs to
24 be identified. In accordance to the present invention, the events might be identified as
25 errors, warnings, and messages. The values listed below event indicate the type of error,
26 warning, or message that occurred within the functional area or component being traced.
27

The last column allows for robustness within the present invention. The user is able to add other variables, functions, etc. that are desired to be traced. For example, as the maturity of a product increases, it may become necessary to record more information. Additional variables may be added under the free1 column to allow for this added storing of information. Of course, the errors, messages, and warnings may be stored within separate buffers.

Table 2 is an example of a buffer storing the tracing of a second component of the computer program where the results of tracing the components are stored (e.g., 412, 416, and 420) in separate post-processing buffers (e.g., buffers 214 of Figure 2).

Component 2 Trace

Errors

Time	Function	State1	State2	Vars1	Vars2	Event	Free1
25553010	FunctionName39	9B20	0000	00FD	0005	2046	0000

Warnings

Time	Function	State1	State2	Vars1	Vars2	Event	Free1
25341009	FunctionName48	9B20	0000	000F	0151	204B	FAA8
25784420	FunctionName46	9B20	0000	000A	0005	204C	5BCD

Messages

Time	Function	State1	State2	Vars1	Vars2	Event	Free1
28600716	FunctionName48	9B20	0000	000F	0151	2001	FAA8
29124801	FunctionName50	9B20	0000	000D	0005	2002	0000
29984775	FunctionName50	9B20	0000	000D	0005	2010	0FFF

Table 2

The definition and description of the columns listed in table 2 are the same as those listed above describing table 1. Similar buffers may likewise be used and filled for each component of a program in this manner.

Once the tracing facility has traced (406) the program code of each component comprising the system, an analysis of the discovery of any types of events occurs. Should there be no events stored in the buffers, the system continues to operate normally. However, if events have been identified, time stamped, and stored in buffers, a merging of events (e.g., step 313) may occur. A request to merge may be initiated by an authorized user through an interface. All events are preferably time stamped (410) with equal granularity and from the same clock. The timestamps may then be used to merge multiple buffers during post processing, enabling a visual representation of when events actually occurred to be analyzed.

Table 3 is one example of the post-processed buffers from Component 1 and Component 2 merged according to the time stamp of each event stored.

Component 1 and Component 2 Merged Trace

Time	Function	State1	State2	Vars1	Vars2	Event	Free1	EWM	Comp.
21684712	FunctionName4	9B24	0000	0005	0151	010F	0000	Message	1
22734407	FunctionName2	9B24	0000	0001	0005	01F5	0000	Message	1
23334000	FunctionName1	9B24	0000	0000	000F	01F3	0000	Message	1
23785393	FunctionName2	9B24	0000	0001	0005	01F5	0000	Message	1
25341009	FunctionName48	9B20	0000	000F	0151	204B	FACC	Warning	2
25553010	FunctionName39	9B20	0000	00FD	0005	2046	0000	Error	2
25684779	FunctionName1	0000	0000	0005	0151	00AC	0000	Error	1
25784300	FunctionName2	9B24	0000	0001	0005	013F	0000	Error	1
25784420	FunctionName46	9B20	0000	000A	0005	204C	5BCD	Warning	2
25784901	FunctionName2	9B23	0000	0005	0151	013A	0000	Warning	1
25790200	FunctionName3	9B24	0000	0001	0005	0122	0000	Warning	1
25790999	FunctionName3	9B24	0000	0001	0005	0122	0000	Warning	1
29124801	FunctionName50	9B20	0000	000D	0005	2002	0000	Message	2
29984775	FunctionName50	9B20	0000	000D	0005	2010	0FFF	Message	2

Table 3

The descriptions and definitions of the columns in table 3 are again the same as the definitions used to described the columns in table 1 and table 2. The additional columns EWM and Comp. Indicate what kind of event was discovered by identifying it as either an error, warning, or message. The column Comp. identifies in which component of the program coded the event occurred.

Though every possible type of event may occur many times and may be identical in nature to the other events from separate components, each event is uniquely identified by type and time and thus, events are not duplicated anywhere within the system. Each task, process, or component of the embedded system may have its own independent set of trace buffers 214.

After the trace buffers from each individual component of the system are merged (e.g., 313) and the stored events are analyzed (e.g., 326), the errors and possibly warnings are corrected by recoding (328) where necessary. Once the recoding has been conducted, the software code resumes operation, and if errors, warning, and messages are still present, the events return to being analyzed (326) in search of a solution to correct the errors.

The tracing of the program code in each component of the system is preferably always enabled. Should a catastrophic error be discovered, all trace buffers are automatically collected. This collection of all trace buffers may also occur in the event of an authorized user-generated request. Collection of the buffers is preferably directed to non-volatile storage that may be accessed at a later time. After subsequent retrieval from non-volatile storage to an off-board destination, post processing of the buffers occurs. Collection of trace buffers also minimizes the need for re-creation of a catastrophic event should a catastrophic error occur.

By providing the separate buffers 214, the present invention ensures that no component consumes buffer resources needed by other components. Separate buffers 214

1 for error, warnings, and messages further ensures that more commonly occurring events are
2 prohibited from overwriting more critical events.

3 Figure 5 provides a method 500 of merging events stored in separate buffers 412,
4 416, and 420 of the present invention. The method 500 starts 502 and the tracing facility
5 obtains 504 the merge parameters. For instance, the merge parameters may indicate how
6 many buffers will be merged and the locations where the merged buffers will be stored. The
7 merge parameters may be set at the time of coding of the program, or may be set by user
8 input. After the merge parameters have been obtained 504, the buffers that are to be merged
9 are located 506. These may include the three separate buffers (e.g., 412, 416, and 420) as
10 described earlier.

11 The method 500 then merges 508 the contents of these separate buffers into one
12 single buffer. The events stored in separate buffers that are merged together 508 are
13 preferably organized 510 according to the time stamp each has previously received. This
14 allows the user to evaluate each event in the order the event occurred. The contents of the
15 merged buffer are then provided 512 to the user, typically by screen dump or printing. The
16 user is able to then examine the content of the merged events and prioritize the importance of
17 each event found by the tracing facility of the present invention.

18 The present invention may be embodied in other specific forms without departing
19 from its spirit or essential characteristics. The described embodiments are to be considered in
20 all respects only as illustrative and not restrictive. The scope of the invention is, therefore,
21 indicated by the appended claims rather than by the foregoing description. All changes which
22 come within the meaning and range of equivalency of the claims are to be embraced within
23 their scope.

24 What is claimed is:
25
26
27